WHITE PAPER · VERSION 1.0 · JUNE 2026

The AI Governance Execution Gap

Why regulated African fintechs need runtime controls, not just AI policies

AUTHOR

Elvis Tapfumanei

REGION

South Africa and African financial services markets

Executive Summary

AI is already inside everyday fintech work. Employees use it to summarise documents, draft customer responses, analyse bank statements, debug code, and make operational decisions faster. South Africa's regulators have taken notice: in November 2025, the FSCA and the Prudential Authority published the country's first comprehensive joint report on AI in the financial sector, finding that banks lead adoption at 52% and payments institutions follow at 50%.[1]

The governance problem is that most organisations are still treating AI as a policy issue when the risk is increasingly operational.

An AI policy can tell employees not to paste customer data into an external model. A vendor questionnaire can confirm a platform has security controls. A training session can explain responsible AI use. But none of these controls answers the most important operational question:

> *What happens at the exact moment an employee, AI agent, workflow, or system attempts to send sensitive data, regulated context, or decisioning logic into an AI model?*

That moment is where the AI governance execution gap appears. The gap is not the absence of principles. Most regulated organisations already understand privacy, fairness, security, accountability, and responsible AI. The gap is the distance between those principles and the runtime layer where actual AI interactions happen.

In fintech, this gap is especially material. A single AI interaction may involve customer identity data, bank statements, affordability information, credit risk logic, complaints history, collections context, fraud signals, payroll data, internal source code, or confidential commercial information. If that interaction is not governed at runtime, the organisation may not know what data moved, which tool received it, whether the action was approved, or what evidence exists after the fact.

South Africa's Information Regulator reported 1,947 security compromise notifications between April and November 2025 (a 40% increase year-on-year), with enforcement fines issued and further action signalled.[2] The upcoming Conduct of Financial Institutions Bill (COFI) will require automated and technology-driven systems to be confirmed as fit for purpose under new conduct expectations.[3]

This white paper argues that regulated African fintechs need to move beyond AI policies, training, and vendor questionnaires towards a runtime governance model: one that classifies AI interactions, enforces policy decisions, coaches users, escalates high-risk events, and produces audit-ready evidence.

The future of AI governance is not a better document. It is a control plane.

1. The Myth of the AI Policy

An AI policy is necessary but not sufficient. Policy creates language, establishes intent, and gives leadership a basis for accountability. The problem begins when the policy is treated as the control.

A policy does not know when an employee opens ChatGPT or Microsoft Copilot. It does not classify the contents of a prompt. It does not detect payroll data in a spreadsheet summary. It does not block customer identity numbers from leaving the organisation. It does not stop an AI agent from producing an unauthorised customer promise. It does not generate an immutable audit trail. It does not prove what happened during an incident.

> *Policy lives at the intention layer. AI risk increasingly lives at the interaction layer.*

This gap produces what might be called governance theatre: the artefacts exist, but when a regulator, auditor, CISO, or executive asks for proof, the organisation still cannot answer basic questions: Which AI tools are employees using? What sensitive data has entered those tools? Which interactions were blocked, coached, or escalated?

The FSCA and PA's 2025 joint survey confirmed this pattern directly. Of approximately 2,100 survey responses across the South African financial sector, only 10.6% of respondents reported active AI use, yet the regulators described the finding as a 'critical knowledge gap.'[4] For low-risk experimentation, governance theatre may remain invisible. For regulated fintech operating under POPIA, the Banks Act, and the incoming COFI regime, it will not.

2. Why Fintech Exposure Is Different

Fintech organisations handle data and decisions that sit close to money, identity, creditworthiness, fraud, complaints, financial vulnerability, and customer conduct. That changes the risk profile of AI adoption significantly.

A fintech using AI to summarise affordability information, interpret support tickets, analyse bank statements, or draft customer eligibility responses faces operational liability beyond simple data leakage.

AREA	TYPICAL AI EXPOSURE
Customer support	AI-generated responses, complaint summaries, eligibility explanations, refund wording
Credit & affordability	Bank statements, income patterns, affordability reasoning, credit decision context
Fraud & risk	Transaction patterns, risk flags, suspicious behaviour, internal detection logic
KYC & onboarding	Identity documents, verification results, personal information, company records
Collections	Vulnerability indicators, repayment arrangements, arrears history, customer commitments
Engineering	Source code, infrastructure details, logs, credentials, incident summaries
Compliance	Regulatory interpretations, policy documents, monitoring outcomes, audit evidence
HR & payroll	Employee records, salary data, disciplinary information, personal circumstances

Under POPIA section 11, personal information may only be processed if there is a lawful basis. Under section 19, responsible parties must implement appropriate technical and organisational measures to secure personal information. And under section 72, personal information may not be transferred to a third party in a foreign country unless the recipient provides substantially similar protection.[5]

These tools process data on infrastructure that may be located in the United States or European Union. Each such interaction is a potential cross-border transfer. The organisation must be able to establish a lawful basis, assess adequacy of the destination's protections, and document that assessment.[6] POPIA uniquely extends these protections to juristic persons, making the compliance obligation broader than comparable frameworks like the GDPR.[7]

3. The Liability Layer: When AI Becomes Operational

AI becomes materially risky when it moves from advice to action. As adoption grows, AI starts influencing operational decisions and customer outcomes. The FSCA and PA's 2025 report found that operations and IT are the primary areas for traditional AI applications, while sales and marketing lead for Generative AI. That is exactly where customer interactions, conduct obligations, and privacy risks converge.[8]

If a customer support AI agent hallucinates a promise, refund, or eligibility statement, it becomes a conduct, complaints, legal, compliance, and evidence problem. Under Treating Customers Fairly principles, institutions are expected to ensure that customers receive clear and fair communications and are not misled.[9]

The same applies to internal workflows. If an AI-assisted analysis influences a lending decision, collections treatment, fraud review, or customer segmentation exercise, the organisation needs more than a record of the final business action. It needs traceability around the AI-assisted process that contributed to that action.[10]

The incoming COFI Bill reinforces this direction. COFI will require financial institutions to maintain adequate operational capabilities, meet reporting and record-keeping requirements, and ensure that automated and technology-driven systems are fit for purpose.[11] The FSCA Commissioner has emphasised that readiness is an industry-wide responsibility.[12]

4. From Policy to Runtime Control

A runtime governance model treats each AI interaction as a controlled event. The control model should answer five questions before or during the interaction:

QUESTION	SCOPE
Who is making the request?	Identity, role, business unit, device, access context, and authorisation.
What is being sent?	Does the prompt, file, query, or action contain personal information, financial data, or confidential business context?
Where is it going?	Approved internal model, approved vendor, external public AI tool, or unknown destination.
Which policy applies?	Policy should depend on jurisdiction, data type, user role, system, workflow, and risk category.
What should happen now?	The decision may be allow, block, redact, coach, escalate, or log.

This is the shift from policy as a document to policy as an executable control.

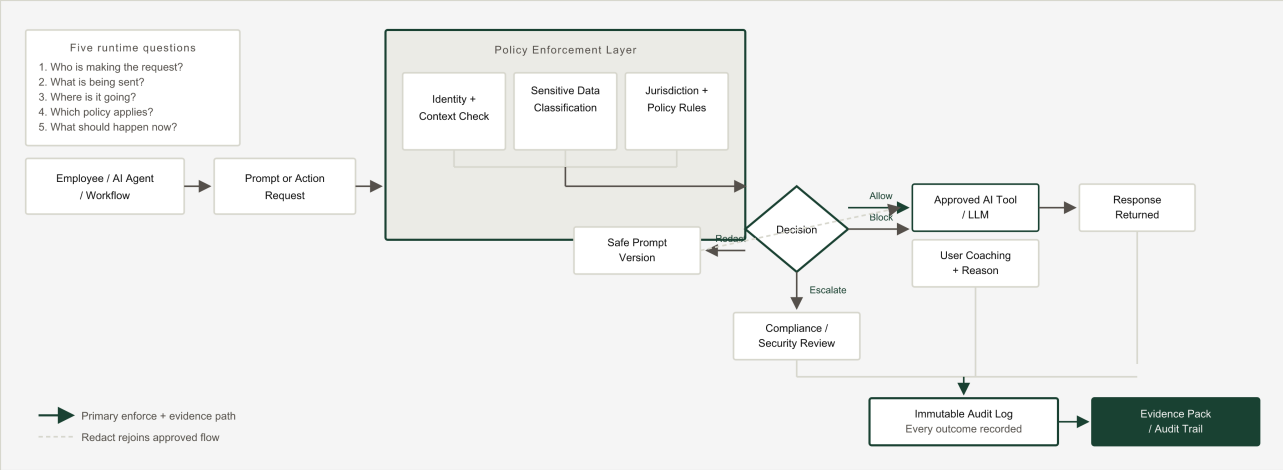

Figure 1: Runtime AI Governance Control Plane diagram

> *The strongest control is not always a hard block. User coaching is often more effective and more operationally sustainable.*

Consider an employee who attempts to paste a customer complaint into an external model. A runtime control may block the prompt and explain why, suggest a safer version with identifiers removed, or redirect the user to an approved internal tool. This matters because governance must support delivery, not simply obstruct it. A governance model that only says no will be bypassed. A governance model that gives clear alternatives becomes part of the operating layer.[13]

5. Evidence Is the Missing Governance Layer

AI governance becomes credible when it can produce proof. For regulated organisations, the relevant question is not 'do we have a policy?' but 'can we prove it operated?' Evidence is the difference between governance as intent and governance as control.[14]

Between April and November 2025, South Africa's Information Regulator received 1,947 security compromise notifications (a 40% increase over 2024). The average cost of a data breach in South Africa in 2024 was R53 million.[15] The Regulator has issued infringement notices and administrative fines, and has signalled that enforcement will intensify.[16]

Evidence Elements

EVIDENCE ELEMENT	WHY IT MATTERS
User identity	Shows who initiated the interaction
Timestamp	Establishes when the event occurred
Tool or destination	Shows where the interaction was directed
Data classification	Shows what type of information was involved
Policy rule applied	Shows which control standard governed the decision
Decision outcome	Allow, block, redact, coach, or escalate
Rationale	Explains why the decision was made
User notification	Shows whether the employee was coached or warned
Tamper-evident record	Strengthens integrity of the evidence
Exportable evidence pack	Allows audit, incident review, or regulatory response

6. The South African Fintech Context

South African fintechs cannot simply import AI governance models designed for other regions. The South African regulatory environment has its own structure and timing.

6.1 POPIA and the Cross-Border Transfer Problem

When employees use external tools, customer data is typically processed outside South Africa. Under POPIA section 72, a responsible party may not transfer personal information to a foreign country unless that recipient provides substantially similar protection.[17] South Africa does not maintain a formal list of adequate countries; the burden falls on each responsible party to assess and document adequacy.[18] Every uncontrolled use potentially constitutes an undocumented cross-border transfer.

6.2 The COFI Bill and Conduct Obligations

The Conduct of Financial Institutions Bill will establish a principles-based market conduct regime. Institutions must demonstrate that their processes and systems achieve desired conduct outcomes.[20] The COFI Bill is expected to be submitted to Cabinet in late 2025 or early 2026, with promulgation anticipated in 2026.[22]

6.3 The Joint Cybersecurity Standard

In June 2025, the Joint Standard on Cybersecurity and Cyber Resilience Requirements came into effect. The standard requires financial institutions to implement governance structures, conduct regular risk assessments of third-party providers, and maintain operational resilience.[23]

7. The Composite Scenario: What an Incident Looks Like

Consider a credit operations analyst who, under time pressure, pastes a spreadsheet of customer bank statement data into ChatGPT. From a governance perspective, four things have happened: personal information has been transferred without an adequacy assessment; financial data has been processed by an unapproved model; the output may influence prioritisation logic; and there is no audit trail.

> *When the incident is discovered, the organisation is in a forensic reconstruction exercise with no evidence and no controls to point to.*

Now consider the same scenario with a runtime governance model. When the analyst attempts to paste the data, the system classifies it as personal financial information and flags the unapproved destination. The user receives in-line coaching: "this data category requires an approved internal tool." The interaction is logged. If the analyst overrides, it escalates to compliance. The risk is contained before the data leaves.

8. Operating Model: Who Owns AI Governance?

FUNCTION	ROLE IN AI GOVERNANCE
Board / Exco	Sets risk appetite; approves framework
CISO / Security	Technical control, data leakage, monitoring
Compliance	Obligation mapping (POPIA, COFI, TCF); evidence review
Legal	Liability, cross-border transfer risk
Privacy	Personal information processing, section 72 controls
Product	Approved use cases and workflows
Engineering	Technical controls, logging, secure architecture
Operations	Frontline adoption and process impact
Delivery Lead	Converts intent into delivery flow and progress

The Delivery Lead role is consistently under-recognised. Without delivery discipline (backlog prioritisation, sprint planning, stakeholder alignment), AI governance remains a committee output rather than an operational capability.

9. AI Governance Maturity Model

LEVEL	NAME	DESCRIPTION	RISK STATE
1	Unaware	Informal AI use with no visibility.	Shadow AI risk
2	Policy-only	Policy exists but enforcement is manual.	Governance theatre
3	Monitored	Usage is logged but not controlled in real time.	Detection only
4	Enforced	Interactions are classified/blocked at runtime.	Active control
5	Evidence-ready	Every interaction produces audit-ready proof.	Defensible

10. Self-Assessment Checklist

Answer yes or no for each question. Use the scoring guide below to estimate your maturity level.

QUESTION	YES	NO
Do you know which AI tools employees are using today?		
Can you distinguish approved AI use from unapproved AI use in real time?		
Can you prevent customer, payroll, or KYC data from entering external AI tools?		
Have you assessed cross-border transfer risk under POPIA section 72?		
Can you prove which AI interactions were allowed, blocked, or coached?		
Can compliance review an AI incident using a complete evidence trail?		
Do your AI policies map to enforceable technical controls?		
Do you have a clear owner for AI runtime governance?		
Can your current controls cover ChatGPT, Copilot, Gemini, and Claude?		
Can you produce regulator-ready evidence within 24 hours?		

> SCORING GUIDE
> **0–3 yes:** Levels 1–2 · **4–6 yes:** Level 3 · **7–8 yes:** Level 4 · **9–10 yes:** Level 5

11. Implementation Roadmap

Regulated fintechs should sequence the work, starting from highest exposure.

Phase 1: Visibility

Create an inventory of AI tools, use cases, teams, and data types. Most organisations will discover adoption is broader than assumed.

Phase 2: Policy Rationalisation

Simplify policy into enforceable rules. Map allowed, restricted, and prohibited use cases. Define data classification rules.

Phase 3: Runtime Control Pilot

Deploy controls around highest-risk interaction points first. Start narrow. Prove the model. Expand.

Phase 4: Evidence and Review

Turn events into reviewable governance evidence. Build tamper-evident logs and compliance dashboards.

Phase 5: Scale and Optimise

Expand coverage to the broader enterprise. Integrate with existing security and GRC tooling.

12. Conclusion: Governance Must Produce Proof

For regulated South African fintechs, AI governance has moved from principle to execution. At this regulatory moment, the relevant question is not whether the organisation has an AI policy, but whether it can prove that policy operated when it mattered.

The organisations that answer these questions well will not merely be more compliant. They will be more operationally resilient. The future of AI governance is a runtime control layer that turns intent into action and proof.

References

1. FSCA and Prudential Authority. Artificial Intelligence in the South African Financial Sector. November 2025.
2. Information Regulator of South Africa. Media Briefing on POPIA Enforcement. November 2025.
3. Protection of Personal Information Act 4 of 2013 (POPIA). Republic of South Africa.
4. CMS Law. Managing Cross-Border Data Transfers. July 2022.
5. ITIF. South Africa's Cross-Border Data Transfer Regulation. June 2025.
6. Recording Law. South Africa Data Privacy Laws: Complete POPIA Guide. 2026.
7. ENS Africa. FSCA and PA Publish Landmark Report on AI. December 2025.
8. Baker McKenzie. South Africa: AI Adoption by the SARB and FSCA. December 2025.
9. Masthead. 2025 FSCA Three-Year Regulation Plan. 2025.
10. African Law & Business. South Africa Considers Revamped Financial Rules. July 2024.
11. Polity.org.za. The COFI Bill: What Financial Institutions Need to Know. August 2025.
12. ENS Africa. Is the COFI Bill Finally Kicking In? 2025.
13. Caveat Legal. What is the COFI Bill? June 2023.
14. Axiomatic. COFI Bill 2025: Reforming South Africa's Financial Sector. January 2025.
15. LexAfrica / Werksmans. Code Red to Code Regulated. January 2026.
16. Corbado. 10 Biggest Data Breaches in South Africa. 2026.
17. ITWeb. InfoReg Exposes POPIA Violators. November 2025.
18. Coetzee, J. Cross-Border Data Flows and POPIA. PER Journal, 2024.
19. Cliffe Dekker Hofmeyr. Overview of the AI in SA Financial Sector Report. December 2025.
20. Masthead. AI in the SA Financial Sector — What FSPs Need to Know. February 2026.

This white paper is published for informational purposes only and does not constitute legal advice. Readers should obtain appropriate legal, compliance, and technical guidance before taking action on any matter discussed herein.

Ready to close the execution gap?

If this resonates, get in touch. I work with regulated fintechs on runtime AI governance, delivery, and advisory.

ELVISTAPFUMANEI.COM